

GI-Dagstuhl Seminar 15504
Formal Evaluation of Critical Infrastructures
December 6-9, 2015

Monday December 07, 2015

09:00-09:30 Welcome and introduction of participants

09:30-10:25 Invited talk
Joost-Pieter Katoen Verifying Markov Models at a Glance

10:25-10:45 Break

10:45-12:00 Applications

Marijn Jongerden Abstract Energy Resilience Modeling for Smart Houses
The use of renewable energy in houses and neighbourhoods is very much governed by national legislation and has recently led to enormous changes in the energy market and poses a serious threat to the stability of the grid at peak production times. One of the approaches towards a more balanced grid is, e.g., taken by the German government by subsidizing local storage for solar power. While the main interest of the energy operator and the government is to balance the grid, thereby ensuring its stability, the main interest of the client is twofold: the total cost for electricity should be as low as possible and the house should be as resilient as possible in the presence of power outages. Using local battery storage can help to overcome the effects of power outages. However, the resulting resilience highly depends on the battery usage strategy employed by the controller, taking into account the state of charge of the battery. We present a Hybrid Petri net model of a house (that is mainly powered by solar energy) with a local storage unit, and analyse the impact of different battery usage strategies on its resilience for different production and consumption patterns. Our analysis shows that there is a direct relationship between resilience and flexibility, since increased resilience, i.e., reserving battery capacity for backup, decreases the flexibility of the storage unit.

Vince Molnár Verification of a Safety-Critical Module of the Paks Nuclear Power Plant
The Paks Nuclear Power Plant provides more than 40% of the power generated in Hungary, being the first and only operating nuclear power station of the country. It is continuously being maintained and modernized in order to raise power output and extend its lifetime, requiring assessment and verification of new parts and technologies. We present a case study where we applied model checking to verify safety, liveness and fairness requirements on a safety logic used to initiate an emergency prevention action in case the so-called PRISE event occurs. In case of a PRImary to SEcondary (PRISE) leakage accident, radioactive water of the primary circle escapes through the heat exchanger and contaminates the water in the secondary circle. A possible countermeasure is to unload the contaminated water. The model of the PRISE logic was given as a functional block diagram. This was translated into a colored Petri net model, and most of the textual requirements were formalized in branching-time (for safety and liveness properties) and linear (for fairness properties) temporal logic. There were a number of requirements that could not be expressed in temporal logic, which demonstrated the limited applicability of these formalisms in real settings. The analysis was performed by the PetriDotNet tool, which is developed in our research group and capable of verifying both branching-time (CTL) and linear (LTL) temporal logic requirements on regular or well-formed colored Petri nets using symbolic model checking.

Daniel Darvas Formal Methods for the Industrial Control Systems at CERN: Dreams and Reality
The particle accelerator complex of CERN (European Organization for Nuclear Research) relies on various control systems, such as cryogenics, ventilation or vacuum systems. These systems are mainly controlled by Programmable Logic Controllers (PLCs) that are specialized, reliable industrial computers. The control systems are critical to CERN's operation; therefore improving their quality is a high priority. Quantitative formal methods could help us e.g. to assess reliability, improve the energy efficiency, or prove correct scaling of the equipment. However, the gap between the available academic methods and the industrial application seems to be important both in terms of usability and scalability. This does not mean that formal methods could not be used for the critical systems of CERN. In my talk I will briefly overview the possible usage of quantitative formal methods. Then I will focus on the formal verification and specification methods we are working on, together with the recent efforts made to provide usable and scalable solutions for CERN.

12:00-13:30 Lunch break

13:30-14:45 Cyber-physical Systems

Enno Ruijters Reliability Centered Maintenance via Statistical Model Checking

The current trend in infrastructural asset management is towards risk-based (a.k.a. reliability centered) maintenance, promising better performance at lower cost. By maintaining crucial components more intensively than less important ones, dependability increases while costs decrease.

This requires good insight into the effects of maintenance on the dependability and associated costs. To gain these insights, we propose a novel framework that integrates fault tree analysis with maintenance. This framework supports a wide range of maintenance procedures and dependability measures, including system reliability, availability, mean time to failure, as well as maintenance and failure costs over time, split into different cost components.

Technically, our framework is realized via statistical model checking, a state-of-the-art tool for flexible modelling and simulation. Our compositional approach is flexible and extensible.

This presentation will explain our framework, and demonstrate it on two case studies from industrial practice: insulated railway joints, and pneumatic compressors for trains.

András Vörös Hierarchical Runtime Verification of Cyber-physical System

Cyber-physical systems are usually too complex to be verified by offline techniques. However, on-line verification, namely runtime verification can still help. In this presentation I will describe our ongoing work on a small "industry-like" case study. The case-study is a safety critical distributed system to ensure the safety of a railway system by observing the situations locally and decide about the correctness globally. However, for such systems, multiple levels of safety would be desirable. In my presentation I will show how a hierarchical runtime verification approach used to further increase safety. Local decisions are verified locally by using the model of the safety logic. In addition, a high level (system level) logic is also developed to observe the global states of the system. The component level monitors provide information for the system level logic to find problems and errors as fast as possible.

In my presentation I will introduce the components we developed and I will describe further directions for developments.

Justyna Chromik Improving the Security of Critical Infrastructures using a Model of the Physical System

Critical infrastructures, such as for the generation and distribution of the electrical energy, depend on the control networks which monitor and control the physical systems. These control networks, often called SCADA (Supervisory Control and Data Acquisition) networks, were not designed with a notion of security in mind. In the increasingly interconnected world, the SCADA networks are highly vulnerable. Although their security is finally getting more attention in research, many approaches use protection techniques applicable for other networks, such as corporate networks. One of the methods used to improve the security of any network is by monitoring the traffic, searching for patterns and anomalies indicating an intrusion in the network. These, so called, intrusion detection systems (IDS) can alarm network operators about a possible security breach. However, control networks and corporate networks are different. Therefore, the majority of patterns and anomalies for corporate networks is not applicable in control networks. An important difference, which has been gaining increased attention in the past decade, is the idea of incorporating a model of the physical system to the IDS. This approach is taken in order to improve the quality of the IDS, i.e. ensuring less errors and more accuracy in detecting intrusions. Comparing the outcome of control commands on the model of the physical system, before executing it in the real system, can help avoiding undesired states of the physical system. The aim of my PhD project is to improve intrusion detection mechanisms in control networks of electricity distribution, by using a model of the physical system. A first approach in doing so will be presented.

14:45-16:45 Break out / working session

Alexander Gouberman Two-time-scale Markov Reward Models with Rate and Impulse Rewards

Stochastic systems that arise in applications are often modelled in terms of Markov chains which have a large underlying state space and are thus computationally intractable. Often, the transitions in such models exhibit two (or more) different speed scales: fast and slow transitions. Such a qualitative separation of speed can be turned into a quantitative separation by introducing a perturbation parameter (ϵ) and mapping the original Markov chain to a two-time-scale Markov chain (a singularly perturbed Markov chain). The asymptotic behavior of the latter as $\epsilon \rightarrow 0$ provides an approximation of the original Markov chain and often leads to a dramatically reduced state space.

In order to evaluate performance measures of such systems, one can equip the original Markov chain with rate rewards and impulse rewards which leads to value functions from accumulation of such rewards over time. The perturbation parameter ϵ leads to perturbed value functions, and their asymptotic behavior as $\epsilon \rightarrow 0$ gives approximations to the original value function. It turns out that rate rewards lead to regularly perturbed value functions with nice asymptotic properties. In contrast, impulse rewards lead to value functions that remain singularly perturbed, making it necessary to choose a suitable space of value functions equipped with a suitable topology which in turn allows to establish approximations.

Daniel Krähmann Ration and Weight Quantiles

Several types of weighted-automata models and formalisms to specify and verify constraints on accumulated weights have been studied in the past. The lack of monotonicity for weight functions with positive and negative values as well as for ratios of the accumulated values of non-negative weight functions renders many verification problems to be undecidable or computationally hard. In this talk I will present polynomial-time algorithms for computing ratio and weight quantiles in Markov chains, which provide optimal bounds guaranteed almost surely or with positive probability on, e.g., cost-utility ratios or the energy conversion efficiency. This is a joint work with Christel Baier, Clemens Dubsclaff, and Jana Schubert.

Tuesday December 08, 2015

09:00-09:55 Invited talk

Anne Remke and Survivability Evaluation of Critical Infrastructures
Boudewijn Haverkort

09:55-10:15 Break

10:15-11:55 Markov Decision Processes

Nils Jansen Safety-Constrained Reinforcement Learning

We consider controller synthesis for stochastic and partially unknown environments in which safety is essential. Specifically, we abstract the problem as a Markov decision process in which the expected performance is measured using a cost function that is unknown prior to run-time exploration of the state space. Standard learning approaches synthesize cost-optimal strategies without guaranteeing safety properties. To remedy this, we first compute safe, permissive strategies. Then, exploration is constrained to these strategies and thereby meets the imposed safety requirements. Exploiting an iterative learning procedure, the resulting policy is safety-constrained and optimal. We show correctness and completeness of the method and discuss the use of several heuristics to increase its scalability. Finally, we demonstrate the applicability by means of a prototype implementation.

Yuliya Butkova Optimal Continuous Time Markov Decisions

In the context of Markov decision processes running in continuous time, one of the most intriguing challenges is the efficient approximation of finite horizon reachability objectives. A multitude of sophisticated model checking algorithms have been proposed for this. However, no proper benchmarking has been performed thus far.

This work presents a novel and yet simple solution: an algorithm, originally developed for a restricted subclass of models and a subclass of schedulers, can be twisted so as to become competitive with the more sophisticated algorithms in full generality. As the second main contribution, we perform a comparative evaluation of the core algorithmic concepts on an extensive set of benchmarks varying over all key parameters: model size, amount of non-determinism, time horizon, and precision.

Florian Niedermeier Model Checking for Data Centers in Demand Response Systems

Data centers provide critical services to our modern society, e.g. communication or health applications. At the same time, they are major consumers of electrical energy and therefore have a significant impact on the state of the power grid. With increasing volatility in power generation (due to a rising amount of renewable generation), providing flexibility in power demand can be used to support grid stability. Running both critical systems to their mutual benefit, that is, adapting data center power demand to reduce stress on the grid and providing data centers with enough power to offer their critical services is a proposed solution. However, algorithms to perform an adaptation of power demand have to be proved to be of high reliability to comply with required functional safety standards. Model checking is a formal method that may provide the required quantitative and qualitative results. Two properties have to be checked: On the one hand, critical services may not be impaired by allocating too little power to the data center. On the other hand, a predictable amount of power adaptation has to be available with high confidence.

Sascha Wunderlich Greener Bits with Probabilistic Model Checking

Today, in a time when the internet is crucial, data centers are a critical part of our infrastructure. They are well suited for decentralized storage and computation. However, they are also hungry for energy.

One possibility for quenching this hunger is to use local renewable energy sources like solar or wind. The advantage of this method is that it is not only relatively friendly for the environment, but also cheap. The disadvantage is that its availability wildly fluctuates depending on the season, time of day and most importantly the weather.

In this setting, it is desirable to adjust the scheduling of jobs in a data center according to the weather forecast, such that the energy demand adapts to the available green energy.

In this talk, we investigate how probabilistic model checking can be used to achieve this goal. Specifically, we will look at the modeling of the given scenario as a weighted Markov decision process and at the formulation of deadlines and other constraints as according specifications. We will then use existential probabilistic model checking and scheduler synthesis to extract an energy-efficient working plan.

12:00-14:00 Lunch break

14:00-16:00 Break out / working session

16:00-17:40 Synthesis

Sebastian Junges Challenges in PROPhESY - Parameter Synthesis in Markov Chains

We review PROPhESY, our tool for analyzing parametric Markov chains. It can compute a rational function (i.e., a fraction of two polynomials in the model parameters) for reachability and expected reward objectives. PROPhESY supports incremental automatic parameter synthesis (using SMT techniques) to determine “safe” and “unsafe” regions of the parameter space. All values in these regions give rise to instantiated MCs satisfying or violating the (conditional) probability or expected reward objective. PROPhESY features a web front-end supporting visualization and user-guided parameter synthesis. Experimental results show that PROPhESY scales to Markov chains with millions of states and several parameters. In this talk, we discuss challenges related to the underlying state elimination algorithm to compute the rational function as well as the selection of suitable candidates for the partitioning of the parameter space.

Arnd Hartmanns Stochastic Hybrid Automata: The Model, the Tools, and the Challenges

Stochastic hybrid automata (SHA) are an expressive unifying model that provides for discrete and continuous nondeterminism, stochastic choices, real-time behaviour, and complex continuous dynamics. They constitute the formal semantics of the “Hybrid Modest” extension of the high-level Modest modelling language. The analysis of SHA is supported by a model-checking technique based on a combination of abstraction, hybrid automata reachability, and probabilistic model checking. It delivers upper/lower bounds on maximum/minimum reachability probabilities and expected rewards, and can thus be used to verify safety. However, its application in practice is hindered by a lack of stable and scalable tool support. This is partly rooted in the technical challenges of adapting existing hybrid automata solvers, but also partly because the reachability abstractions computed by those solvers are not designed for the requirements of the subsequent probabilistic model checking task. In this talk, I will present the model of stochastic hybrid automata, describe the existing analysis technique, and outline the challenges - both technical and conceptual - faced by previous attempts to implement SHA model checking tools.

Chih-Hong Cheng Software Synthesis towards Practice

To bring Linear Temporal Logic (LTL) synthesis from theoretical results to practice, by observing current research advances so far, one unavoidably needs to provide answers to two fundamental challenges - The first is related to the *ease of specification*, where one often starts with concrete implementation in mind but disappointedly writes formal specifications that are not realizable. The second is about *code understandability*, where apart from synthesis speed, state-of-the-art LTL synthesis algorithms always generate implementations that are hardly understandable, while implementations from manual programming are commonly elegant, simple, and understandable by engineers.

Towards this goal, we take a pragmatic approach by benchmarking more than 60 different (i.e., each describes a unique application scenario) Programmable Logic Controller (PLC) training examples offered by industrial automation companies or available online. Then, we formally model the requirements using LTL and continuously revise the specification until the synthesized implementation mimics the actual implementation. The set of benchmarks offers important insights towards the commonly seen problems in a particular domain. By doing so, we derive “true specifications” that one should consider when doing specification analysis. From those specifications, we identify specification “patterns” that are used in functional synthesis of PLC programs. As the number of discovered patterns are limited to a few, by importing templates based on the identified patterns to the specification language, one largely mediates the problem induced by the hardness of using LTL as the specification language. More surprisingly, the pattern-based approach naturally derives a pragmatic resolution-based method which not only synthesizes controllers with simpler structure but also substantially faster in terms of synthesis speed.

Alexandros Nikou Cooperative Task Planning Synthesis for Multi-Agent Systems Under Timed Temporal Specifications

Cooperative control of multi-agent systems has traditionally focused on designing local control laws in order to achieve tasks such as consensus, formation, network connectivity, and collision avoidance. Over the last decade or so, the field of control of multi-agent systems with complicated behavior under complex high-level task specifications has been gaining significant research attention. In this context, we consider the framework in which each agent is associated with a set of tasks, given in Metric Interval Temporal Logic (MITL) formula, that should be fulfilled both from each individual agent as well as the team of agents. We propose a method for automatic control synthesis in a two stage systematic procedure. With this method we guarantee that all the agents satisfy their own individual task specifications as well as that the team satisfies a single team global task specification. The solution involves a sequence of Automata constructions.

20:00 ROCKS board meeting

Wednesday December 09, 2015

09:00-09:55 Invited talk

Verena Wolf Aggregation of Stochastic Models of Complex Networks

09:55-10:15 Break

10:15-10:55 LTL with and without Büchi

Yong Li Model Checking Fairness LTL Efficiently

We propose a new algorithm to verify fairness LTL formulas. We first present an efficient algorithm to deal with a fragment of fairness formulas and then extend the algorithm to handle arbitrary ones. Notably, by making use of some syntactic transformations, our algorithm avoids to construct corresponding Buchi automata for the whole fairness formulas, which can be very large in practice. We implement our algorithm in NuSMV and consider a large selection of formulas. Our experiments show that in many cases our approach exceeds the automata-theoretic approach up to several orders of magnitude, in both time and memory.

David Müller

Almost Universality for Unambiguous Büchi Automata

The growing complexity and dependence on computational systems in our every day life requires formal methods for verification to ensure liveness and safety. Probabilistic model checking determines the likelihood that a given system model, e.g. a Markov chain, satisfies a given property.

In this talk we will revisit the problem of automata-based Markov chain analysis. In the literature Markov chain analysis for separated Bchi automata has been covered by Couvreur et al. They show that computing the probability for a Markov chain to generate a trace accepted by the separated Bchi automaton can be done in polynomial time. We will extend the class of separated Bchi automata to unambiguous Bchi automata. Our focus is on the case where the underlying probability measure is induced by the uniform distributions for finite words of the same length. We will present a polynomially time-bounded algorithm for solving the almost universality problem “is $Pr(L(U)) = 1$?”.

For this we will rely on the construction of so called cuts. If we set the initial states of an UBA to one of its cuts, then the UBA becomes almost universal. The existence of a cut yields the probabilistic non-emptiness of the automaton, i.e. $Pr(L(U)) > 0$.

11:00-12:00 Discussion

12:00 Lunch
